

Probing as Quantifying Inductive Bias

Alexander Immer*,¹ Lucas Torroba Hennigen*, Vincent Fortuin, Ryan Cotterell
ETH Zürich MIT
alexander.immer@inf.ethz.ch lucastor@mit.edu
fortuin@inf.ethz.ch ryan.cotterell@inf.ethz.ch

Abstract

Pre-trained contextual representations have led to dramatic performance improvements on a range of downstream tasks. Such performance improvements have motivated researchers to quantify and understand the linguistic information encoded in these representations. In general, researchers quantify the amount of linguistic information through probing, an endeavor which consists of training a supervised model to predict a linguistic property directly from the contextual representations. Unfortunately, this definition of probing has been subject to extensive criticism in the literature, and has been observed to lead to paradoxical and counter-intuitive results. In the theoretical portion of this paper, we take the position that the goal of probing ought to be measuring the amount of inductive bias that the representations encode on a specific task. We further describe a Bayesian framework that operationalizes this goal and allows us to quantify the representations' inductive bias. In the empirical portion of the paper, we apply our framework to a variety of NLP tasks. Our results suggest that our proposed framework alleviates many previous problems found in probing. Moreover, we are able to offer concrete evidence that—for some tasks—fastText can offer a better inductive bias than BERT.¹

1 Introduction

Improved pre-trained representations have led to new performance heights on NLP applications. This has prompted researchers to analyze these representations in an attempt to determine which linguistic properties they encode. **Probing** is the primary method to perform such a quantification; typically, probing consists of training a supervised model, called a **probe**, to predict a linguistic property directly from the representations. It has been

argued that the existence of a high-performing probe suggests that the representation encodes the property of interest (Alain and Bengio, 2017; Belinkov and Glass, 2019). However, despite the apparent simplicity of probing and its wide-spread use, the community has yet to find consensus on several important problems about the endeavor. We enumerate several problems with the supervised probing framework in the following paragraphs.

Problem I (Representation Selection). Counter-intuitively, probing may fail to capture observed differences between representations. For instance, in some supervised probing studies, researchers have shown that random representations are equally good or better than trained ones (Zhang and Bowman, 2018; Pimentel et al., 2020a). This is certainly a nonsensical result; random representations, by construction, do not encode any linguistic property.

Problem II (Probe Selection). There is an ongoing debate on the choice of probes: initially, linear probes were proposed to test the linear separability of learned representations (Montavon et al., 2011; Alain and Bengio, 2017; Liu et al., 2019a). However, more recently, neural networks have been applied with the explicit goal of extracting as much information as possible from the representations (Adi et al., 2017; Conneau et al., 2018; Pimentel et al., 2020b; Pimentel and Cotterell, 2021). Not surprisingly, it has been found that more complex probing tasks often require more complex probes (Belinkov and Glass, 2019). To reduce the risk of overfitting, recent methods aim at trading off probing performance with the probe's complexity (Hewitt and Liang, 2019; Pimentel et al., 2020a; Voita and Titov, 2020).

Problem III (Task Selection). The relationship between probing tasks and NLP tasks remains unclear. This lack of clarity manifests itself in several ways. Firstly, while some argue that probing should

*Equal contribution.

¹Our code is available at <https://github.com/rycolab/evidence-probing>.

focus on simple tasks (Conneau et al., 2018), others argue that probing should focus on complex tasks to be informative (Pimentel et al., 2020a). Thus, it is unclear where to place the boundary between probing and regular NLP tasks and whether there should even be a distinction between the two types of tasks at all. Secondly, how researchers should interpret experimental probing results is still up for debate. For instance, knowing that BERT excels at text generation, is it really surprising that we can predict the tense of a word from a BERT representation? Indeed, the NLP community is still in search of how probing can be of service to downstream tasks.

This paper proposes a new framework for supervised probing that seeks to address the problems described above. We propose to compare representations in terms of the **inductive bias** they provide for a particular task. This may seem counterintuitive, since classical machine learning often refers to the inductive biases of models alone, and not of representations; however, we propose to instead think of models as representation–probe pairs. Such a paired model takes raw text as input, converts it into a representation, e.g., using BERT (Devlin et al., 2019), and predicts a property of interest using a probe. We formalize the notion of the inductive bias of a paired model using the Bayesian model **evidence**. The evidence naturally trades off performance and complexity (Rasmussen and Ghahramani, 2000; MacKay, 2003; Bishop, 2006), therefore, it is well-suited to quantify the amount of inductive bias that a representation–probe pair provides for a particular task.

Indeed, we argue that, by quantifying inductive biases using the evidence, we can solve the problems listed above. The evidence inherently penalizes random representations, addressing **Problem I**, and allows us to automatically select probes that have the right complexity for the given task and representation, addressing **Problem II**. Importantly, automatically controlling probe complexity leads to an apples-to-apples comparison among representations, since every representation has access to the probe best suited for it. For example, we now have a fair basis for comparison between acontextual fastText representations and contextual BERT representations. Finally, evidence-based probing unifies probing and task-driven NLP (**Problem III**): the goal of the experimenter should be to identify the representation–probe pair

with the best inductive bias for a particular problem so there is no difference in how the framework handles probing tasks and regular NLP tasks.

To validate our framework, we apply it to 28 tasks, many of which have been used for probing before. Our results suggest that our framework provides a practical solution to **Problem I** and **Problem II**. With respect to **Problem I**, we never find that random representations encode more inductive bias for a task than pre-trained representations. With respect to **Problem II**, we find that the optimal choice of probe depends on the task and representation in question, e.g., when relying on random representations, a linear probe suffices (since the added complexity of a neural probe cannot possibly help); however, with BERT representations, sometimes it is better to use a non-linear probe. This suggests that our method automatically gets around the probe selection problem. Moreover, our results also suggest that fastText can provide a better inductive bias than BERT for some morphosyntactic probing tasks.

2 Probing as Quantifying Inductive Bias

At the most fundamental level, the NLP community’s interest in pre-trained representations is about reducing the sample complexity of models on downstream tasks. The community hopes that pre-trained representations are able to imbue NLP models with enough information about a given language that models can reach a higher performance with the same or even fewer training data. And, indeed, over and over again this has been shown to be the case (Peters et al., 2018; Devlin et al., 2019; Raffel et al., 2020). Another way of phrasing this desire is that the NLP community hopes that pre-trained representations have a suitable inductive bias for downstream tasks. This paper takes the position that, rather than probing the pre-trained representations for how much linguistic structure they contain—an endeavor that has received much attention (Belinkov et al., 2017; Belinkov and Glass, 2019; Conneau et al., 2018; Liu et al., 2019a, *inter alia*) but is still contentious (Hewitt and Liang, 2019; Pimentel et al., 2020a,b; Voita and Titov, 2020)—we should directly ask how much they improve the inductive bias on tasks of interest.

We propose to quantify the inductive bias of a model, i.e., a representation–probe pair, using the principle of **Occam’s razor** (Rasmussen and Ghahramani, 2000). Occam’s razor states that we

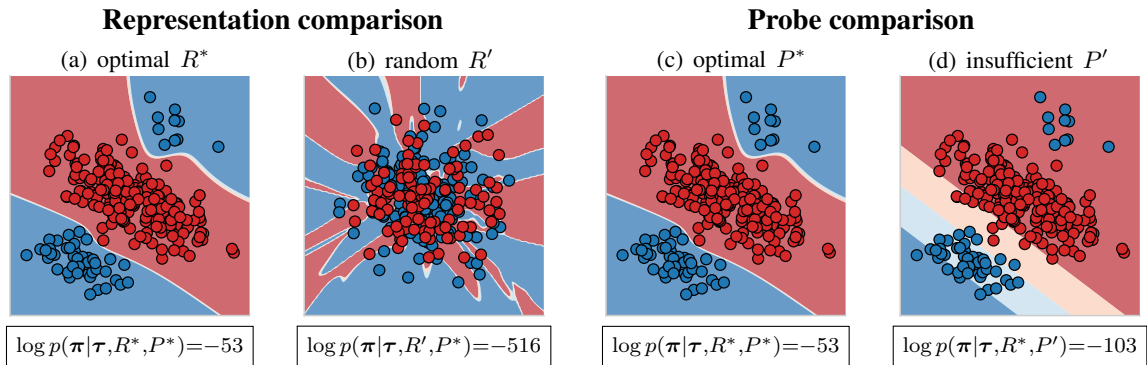


Figure 1: Comparison of the inductive biases of representation–probe pairs using the evidence. The evidence below the respective figures indicates that the right probe and representation are selected. The probing task is a binary classification of two properties (● vs ●). The same colors are used to mark the probe’s decision function. Representations that naturally separate the properties are preferred over random representations in terms of the evidence, since they have a better inductive bias. **Left:** we compare an optimal representation that distinguishes both property classes (a) and a random representation (b). **Right:** we compare a neural probe (c) to a linear probe (d) which is too simplistic. The evidence correctly prefers a neural probe since it better explains the data.

should choose the simplest model that sufficiently explains our observations. One way to operationalize this principle is Bayesian model selection (Rasmussen and Ghahramani, 2000; MacKay, 2003; Bishop, 2006). Bayesian model selection relies on the **evidence**, which is a distribution over data sets for a given model—that is, how likely is it that a particular data set could have been generated by that model. With a probing data set, the evidence encompasses Occam’s razor because (i) a model that is too simple would assign low probability to the data set (e.g., it is very unlikely that we sample a smooth cubic curve from a linear model), and (ii) an overly complex model would assign low probability because it can model that data set as well as many others (e.g., it is unlikely that we sample a cubic from a deep Transformer). In line with Occam’s razor, the evidence is then highest for the simplest model that sufficiently explains the data set (e.g., a cubic model is the best explanation for a data set consisting of a cubic polynomial).

In the following, we outline the probabilistic model for probing and the form of the evidence. This enables us to quantify the inductive bias of representations. Crucially, part of the inference is to select the optimal probe for each representation so as to enable a fair comparison between representations.

2.1 A Probabilistic Model of Probing

Computation of the evidence requires the definition of a probabilistic probing framework. In this section, we introduce such a framework. Specifically, we compute the evidence of representation–probe

pairs that constitute models for a fixed task.²

We start by introducing the notation necessary to describe our probabilistic probing framework. Formally, we denote linguistic sequences by $\tau \in \mathcal{V}^+$, where \mathcal{V} is a vocabulary.³ For example, τ could be a word in context, a whole sentence, or simply a single token. We probe for a linguistic property $\pi \in \Pi$. In a probing task, we have a data set of N i.i.d. pairs $\{(\tau_n, \pi_n)\}_{n=1}^N$ of sequences with associated linguistic properties. We abbreviate all sequences and properties collectively in a data set by τ and π . Formally, a **representation** $R(\cdot)$ is a (possibly stochastic) function from a sequence to a D -dimensional real vector, i.e., $R : \mathcal{V}^+ \rightarrow \mathbb{R}^D$. We will use the shorthand $h = R(\tau)$ to represent the vector resulting from the application of the function $R(\cdot)$ to τ , and h to abbreviate the representations of all sequences τ in the data set. Finally, we employ a probe to predict the linguistic property π_n of a sequence τ_n from its representation $R(\tau_n)$, i.e., a probabilistic probe $f(\cdot)$ maps a vector in \mathbb{R}^D to a distribution over linguistic properties. In all, this means that the composition $(f \circ R)(\tau_n)$ yields a distribution over the linguistic property π_n corresponding to τ_n . As an example, the representation $R(\cdot)$ may be realized by BERT, the probe $f(\cdot)$ may be a linear classifier, τ are words in context, and π are POS tags.

In our framework, we treat the composition of $f(\cdot)$ and $R(\cdot)$ jointly as a single model whose in-

²We note that our formulation has a close connection to the MDL formulation of probing (Voita and Titov, 2020).

³ \mathcal{V}^+ is the set of all sequences of elements in \mathcal{V} of length at least 1.

ductive bias we seek to assess. Formally, we define a model as a **representation–probe pair**, which we denote by a tuple $(R, P) \in \mathcal{R} \times \mathcal{P}$, where $R(\cdot) \in \mathcal{R}$ denotes a representation and $P \in \mathcal{P}$ is a probe specification. A **probe specification** characterizes a prior over some family of probes, e.g., a 2-layer neural network probe with tanh activations and a Gaussian prior on the weights. This is consistent with the probing literature, where probes are often parameterized families of probabilistic models trained using a regularization scheme that implicitly defines a prior over the parameters.⁴ In such a case, a natural prior has the form $p(\theta | \mathbf{h}, P)$, where θ are the parameters of the family of models associated with P .⁵ Each $P \in \mathcal{P}$ would then specify a prior over probe parameters θ and thus probe functions $f(\cdot)$. However, we opt for a slightly different notation. Analogous to our notation for h , we define f for the corresponding vector of probe outputs for an input representation, i.e. $f = f(h)$, and \mathbf{f} as the probe outputs over the entire data set. Then, we reparameterize the prior $p(\theta | \mathbf{h}, P)$ in terms of the probe outputs \mathbf{f} , i.e., $p(\mathbf{f} | \mathbf{h}, P)$.⁶ Our formulation is therefore general: we can follow previous work on probing and opt for a neural network probe, in which case each $P \in \mathcal{P}$ can specify an architecture and prior over parameters; however, we can also consider priors directly on function outputs, e.g., if we want a Gaussian process probe.

As we mentioned above, we allow for stochastic representations $R(\cdot)$. We can interpret this as a prior over representation outputs h , which is given by $p(h | \tau, R)$: it is conditional on the choice of representation and the particular input sequence τ we want a representation for. Formulating representations as probabilistic allows our framework to be more general, i.e., it can be used to compare stochastic representations (Vilnis and McCallum, 2015; Barkan, 2017; Xue et al., 2021, *inter alia*) to deterministic representations like BERT. If $R(\cdot)$ prescribes a deterministic representation then the distribution on h given a sequence τ is given by the

⁴For example, L2 regularization can be seen as placing a Gaussian prior over the parameters of a model (Murphy, 2012, Chapter 7).

⁵In most applications, we would assume that the prior does not depend on \mathbf{h} , i.e., the prior would simply be $p(\theta | P)$. Indeed, we are being more general than what is usually necessary; however, as will soon become clear, allowing for this conditioning will simplify notation.

⁶This reparameterization may be achieved with a standard application of the change-of-variable formula using the neural network’s Jacobian, similar to what is being done in functional variational inference (e.g., D’Angelo and Fortuin, 2021).

Dirac delta function: $p(h | \tau, R) = \delta(R(\tau) - h)$.

Jointly, the priors over probe and representations outputs specify the prior for a representation–probe pair. All that remains is specifying the likelihood function; it is defined such that it factorizes over the data set as $p(\boldsymbol{\pi} | \mathbf{f}) = \prod_{n=1}^N p(\pi_n | f_n)$. The joint distribution $p(\boldsymbol{\pi}, \mathbf{f}, \mathbf{h} | \tau, R, P)$ of the probabilistic probing model is then given by

$$\underbrace{p(\boldsymbol{\pi} | \mathbf{f})}_{\text{likelihood function}} \times \underbrace{p(\mathbf{f} | \mathbf{h}, P)p(\mathbf{h} | \tau, R)}_{\text{prior}}. \quad (1)$$

We obtain the evidence for our representation–probe tuple by integration:

$$p(\boldsymbol{\pi} | \tau, R, P) = \iint p(\boldsymbol{\pi}, \mathbf{f}, \mathbf{h} | \tau, R, P) d\mathbf{f} d\mathbf{h}. \quad (2)$$

The evidence is a distribution over linguistic properties $\boldsymbol{\pi}$ given input tokens τ and a particular choice of model, i.e., representation–probe pair (R, P) . A representation–probe pair that could easily generate correct linguistic properties will score a higher evidence than one that does not generate any linguistically meaningful properties or one that can generate all sorts of data sets.

2.2 Maximizing the Model Evidence

To find the best representation–probe pair, we need to find the one maximizing the evidence in eq. (2):

$$(R^*, P^*) = \operatorname{argmax}_{(R, P) \in \mathcal{R} \times \mathcal{P}} p(\boldsymbol{\pi} | \tau, R, P). \quad (3)$$

The space of representations \mathcal{R} that we compare when probing is typically quite small and leads to a discrete choice: each $R(\cdot) \in \mathcal{R}$ simply denotes a distinct choice of representation. Further, all prior work on probing considers exclusively deterministic representations which, as mentioned above, simplifies the prior over representations to a Dirac delta distribution. This means we can rewrite eq. (2) as follows

$$\begin{aligned} & \iint p(\boldsymbol{\pi}, \mathbf{f} | \mathbf{h}, P) d\mathbf{f} \delta(R(\tau) - \mathbf{h}) d\mathbf{h} \\ &= \int p(\boldsymbol{\pi}, \mathbf{f} | \bar{\mathbf{h}}_R, P) d\mathbf{f}. \end{aligned} \quad (4)$$

where we use $\bar{\mathbf{h}}_R = R(\tau)$ to emphasize that this is the non-random representation of τ according to $R(\cdot)$. This characterizes our probing procedure: we compute this integral independently for each

representation $R \in \mathcal{R}$ and hence the problem in eq. (3) reduces to selecting, for each representation, the probe specification $P \in \mathcal{P}$ that maximizes the evidence. The **inductive bias** of a representation R is the resulting optimal evidence across probes: $\max_{P \in \mathcal{P}} p(\boldsymbol{\pi} \mid \bar{\mathbf{h}}_R, P)$. This procedure can also be understood as hypothesis testing with a likelihood-ratio test (see App. A).

While \mathcal{R} is simply the set of representations that we want to probe, the set \mathcal{P} that characterizes priors on probes is more complex. It is typically a combination of discrete and continuous choices: For example, the number of layers in a neural probe is discrete, but the setting of weight decay is continuous. Moreover, to ensure that the evidence is not limited by a restricted choice of probe architectures, the set \mathcal{P} needs to encompass sufficiently simple and complex probes at the same time. Hence, we construct our prior on probes by incorporating commonly used probes into it: we consider linear (Alain and Bengio, 2017; Adi et al., 2017; Hewitt and Liang, 2019; Liu et al., 2019a; Pimentel et al., 2020a) and more complex neural probes (Pimentel et al., 2020b; Voita and Titov, 2020) paired with weight decay to control complexity (Hewitt and Liang, 2019; Pimentel et al., 2020a). Probing based on a family of probes instead of a fixed architecture is a key difference to other probing frameworks. In fact, in our experiments (§4) we find that different representations perform best with different probe architectures and hyperparameters. This suggests that limiting probing to a single probe configuration might be misleading.

In practice, to maximize the evidence for each representation over \mathcal{P} , we follow the evidence framework by MacKay (1995, 2003) using the scalable implementation proposed by Immer et al. (2021a). This enables us to quantify the inductive bias of a representation (eq. (4)) and maximize it over $P \in \mathcal{P}$ as required by eq. (3), i.e., for each representation we select $\max_{P \in \mathcal{P}} p(\boldsymbol{\pi} \mid \bar{\mathbf{h}}_R, P)$. It also allows us to maximize the integral over a set of infinitely many choices of weight decay strength, to further control the complexity of the probes. As shown in §4, this leads to highly consistent results and alleviates overfitting, which is a problem that even simple linear probes have.

3 Tackling Probing with Evidence

As outlined in §1, current work in probing faces a series of problems. Here we discuss how these

problems are directly addressed by the evidence.

3.1 Problem I (Representation Selection)

Clearly, random representations have no suitable inductive bias for linguistic tasks. Nonsensical results, such as that random representations outperform pre-trained ones (Zhang and Bowman, 2018; Hewitt and Liang, 2019; Pimentel et al., 2020a) simply indicate overfitting, which is strictly penalized in our framework. Compared to pre-trained representations, random representations have low evidence for linguistic tasks because there is no probe that can reliably predict the properties. In Fig. 1a vs. 1b, we illustrate how a random representation is penalized by the evidence. As we will see in §4, our framework consistently assigns lower evidence to the random representations compared to the pre-trained ones.

3.2 Problem II (Probe Selection)

Current probing results are inextricably bound to the choice of probe, yet for probing to provide us with insights about representations, we must break this dependence. For example, one salient issue in probing is that, while pervasive in the literature, there is a spurious association between linear probes and ease of extraction. This is illustrated in Fig. 1, where we can see a linear probe (Fig. 1d) that offers less ease of extraction than a neural probe (Fig. 1c), as measured by the evidence. This means that could obtain misleading results if we restricted our analysis to linear probes. Conversely, we will later see that linear probes can be too complex for some probing tasks and overfit, though the evidence overcomes this problem (Fig. 4). We avoid the problem of selecting a fixed probe by instead choosing a sufficiently large set \mathcal{P} of priors of families of probes and finding the optimal probe specification, within that family, for each representation; as we will see later, the optimal probe varies considerably across tasks and representations. Instead of heuristic arguments about which probe to choose, the evidence provides a statistically sound way to select one in line with a likelihood-ratio test (Neyman and Pearson, 1933).⁷

3.3 Problem III (Task Selection)

In our opinion, an important issue with probing is that the research program has unclear goals. Like much of task-driven NLP, probing is essen-

⁷Refer to App. A for more details.

tially supervised learning with pre-trained representations. We argue that the goal of quantifying and, in particular, maximizing the inductive bias of representation–probe pairs aligns probing with regular NLP: In both cases, one searches for an optimal model at the lowest possible complexity—it does not matter whether the task of interest is simple or complex.

4 Experimental Setup

We evaluate our framework on a series of token, arc, and sentence tasks. Our token- and arc-level tasks are multilingual,⁸ whereas our sentence tasks only consider English. We remove any property values that have less than 20 examples in any of the splits. All our probes are trained using the Adam (Kingma and Ba, 2015) optimizer. For details on hyperparameters, see App. B.

Token-level tasks. For our token-level probing tasks, we probe for part-of-speech (POS) tags, tense, number, and case. We use the setup in Torroba Hennigen et al. (2020), which consists of mapping the UD v2.5 (Zeman et al., 2019) treebanks to the UniMorph schema (Kirov et al., 2018) using the converter by McCarthy et al. (2018), and extracting examples of tokens tagged for the relevant properties. Next, we obtain the representations for each of those tokens in their sentential context (Torroba Hennigen et al., 2020). Finally, we split the resulting vocabulary using a 65–35 train–test split, such that no word appears in multiple splits. While the evidence does not require such a split, we use the split to validate results (cf. Fig. 4).

Arc-level tasks. For our arc-level tasks, we conduct dependency arc labeling (DAL). This consists of classifying the label for a dependency relation given only the representations for the head and dependent of that relation. These are extracted from the UD v2.5 treebanks using the approach in Pimentel et al. (2020a). We use the default UD splits.

Sentence-level tasks. For our sentence-level tasks, we consider four tasks. The first is MultiNLI (Williams et al., 2018), a natural language inference task. The other three are the BoolQ (Clark et al., 2019), Commitment Bank (De Marneffe et al., 2019), and recognizing textual entailment (RTE; Dagan et al., 2006;

⁸We consider a small but typologically diverse set of languages: English (eng), Arabic (ara), Turkish (tur), Marathi (mar), German (deu), and Chinese (zho).

Bar Haim et al., 2006; Giampiccolo et al., 2007; Bentivogli et al., 2009) tasks, which are part of the SuperGLUE benchmark (Wang et al., 2019). If a task requires one or more passages as input, we first obtain a passage-level representations by averaging over all of its tokens.

Representations. In our token and arc tasks, we compare four different representations $R \in \mathcal{R}$: (i) m-BERT (Devlin et al., 2019), (ii) fastText (Bojanowski et al., 2017; Grave et al., 2018), (iii) a random representation (Rand.), which offers no information, drawn i.i.d. from a Gaussian distribution with zero mean and unit variance and the same dimensionality as BERT for each data point, and (iv) a representation that assigns a unique random vector to every word in our vocabulary, so the only information it provides is the identity of the word (Word Ident.). The dimensionality of (iii) and (iv) is the same as that of the BERT representation. For the sentence tasks, we consider (i) Random, (ii) fastText, (iii) BERT, (iv) ALBERT (Lan et al., 2020), (v) RoBERTa (Liu et al., 2019b), (vi) XLNet (Yang et al., 2019), and (vii) T5 (Raffel et al., 2020). App. C lists details on the exact models and implementations used.

Probe Family. In order to ensure fair comparisons, our framework requires us to define a suitably expressive set of priors \mathcal{P} over probe families. In line with most of the probing literature, this includes linear and neural probes with 1 or 2 hidden layers, 100 hidden units, tanh activation, and varying weight decay parameter.

5 Results

We find that our formulation of probing alleviates the problems that we identified in §3. Firstly, the evidence suggests that random representations have an unsuitable inductive bias for linguistic tasks, which is in line with hypotheses from previous research (Zhang and Bowman, 2018; Pimentel et al., 2020a). Secondly, the automatic selection of the right probe architecture using the evidence shows that linear probes are seldom preferred, at least in our token- and arc-level experiments. That said, we also find evidence that even linear probes can overfit, and that the optimal linear probes may require many of their weights to be regularized to zero. Clearly, allowing different probe architectures between representations is beneficial for a fair comparison: simpler representations can profit

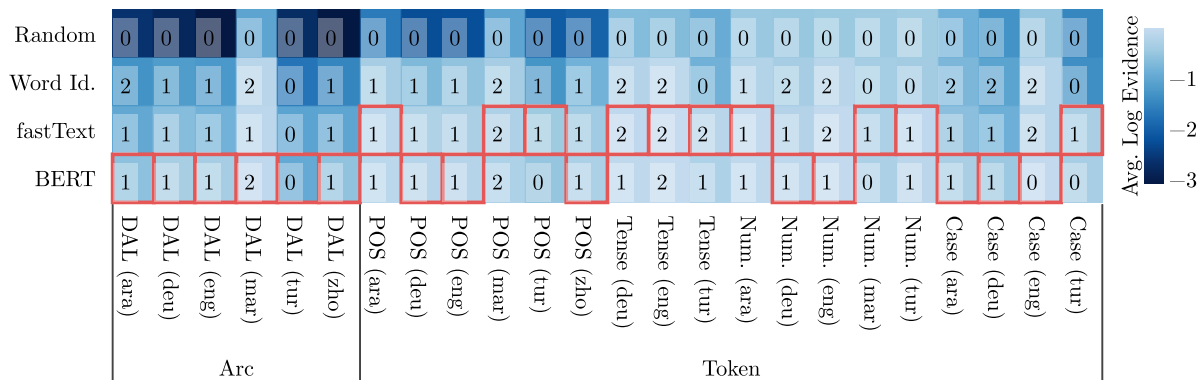


Figure 2: Inductive biases of different representations (shown on the rows) for token- and arc-level tasks (shown on the columns), as measured by the log evidence averaged over the data set. The integers inside the cells denote the number of layers of the optimal probe, with 0 denoting a linear probe. The representations with the best inductive bias (averaged over 3 runs) for each task are highlighted in red.




Figure 3: Inductive biases of different representations (shown on the rows) for sentence-level tasks (shown on the columns), as measured by the log evidence averaged over the data set. The integers inside the cells denote the number of layers of the optimal probe, with 0 denoting a linear probe. The representations with the best inductive bias for each task are highlighted in red.

from a more complex probe and demonstrate a superior inductive bias than more complex representations in some cases. Specifically, we find that fastText demonstrates a better inductive bias than BERT on multiple morphosyntactic tasks, while T5 appears to offer the best inductive bias for all our sentence-level tasks.

5.1 Representation Comparison

In the following, we discuss the results presented in Fig. 2 and Fig. 3 in detail.

Expected trends. Our results depict trends that should be expected from probing. For example, random representations perform worse than pre-trained representations, especially in tasks with a larger number of classes, such as POS and dependency arc labeling. Word identity representations

are better than random representations, which is to be expected, since the former are at least able to associate certain types to their most frequent properties, whereas the latter offer no information because they are sampled randomly per token. We suspect this is the reason why the optimal probe for random representations is always a linear probe that predicts the majority class.

Token- and arc-level tasks. Fig. 2 contains the results of our token- and arc-level tagging tasks. We find that fastText offers a better inductive bias for tense, while BERT is superior for case across all languages with the exception of Turkish (tur). In fact, we find that fastText evinces a better inductive bias for all Turkish token-level tasks. We believe that this is due to the agglutinative nature of Turkish, which means that fastText’s bag-of-subword-units mechanism provides a useful inductive bias. For dependency arc labeling (DAL), we find that BERT has a uniformly better inductive bias. Interestingly, other than for random representations, the optimal probe usually has a non-linearity, which refutes the idea that linear probes should be blindly picked for their simplicity. In all, our token- and arc-level results suggest that BERT is not a panacea, and motivate further research into multilingual studies of the morphosyntactic properties that BERT exposes well.

Sentence-level task. Fig. 3 suggests that T5 (Raffel et al., 2020) has a better inductive bias than the other representations we consider on sentence-level tasks. That said, we find that the difference in evidence between the different representations is generally quite small for BoolQ, RTE, and CB. Indeed, despite these being highly complex tasks, a linear probe is uniformly preferred for BoolQ and RTE.




Figure 4: Average log evidence and cross entropy (*higher is better*) versus weight decay of a linear probe on POS German tasks using BERT. Even a simple linear probe can overfit and yield misleading results if not regularized properly as we can see from the generalization gap between training and test cross-entropies.

This may be an indication that the sentence-level representation mechanism we chose, i.e., averaging over the representations for the tokens in a sentence, is particularly ineffective for these two tasks. Indeed, we see that for both tasks, the evidence for the representations is not much higher than the evidence for the random representation, which may indicate that the optimal probes are largely ignoring the representations and just learning a majority-class baseline, which is achieved at the smallest complexity using a linear probe.

5.2 Controlling Probe Complexity

Fig. 4 shows linear probes on two tasks and how the evidence and cross-entropy change as a function of their weight decay. The graph shows that insufficient regularization leads to poor generalization using BERT, apparent from the gap between training and test loss that grows larger when weak regularization is applied. This means that insufficiently regularizing linear probes—and hence allowing them to fully use their parameters—reduces their evidence.

This observation, alongside former results, led us to conjecture that optimal probes may actually be restricted linear models, i.e., linear probes where most parameters are disabled. Our implementation is easily able to account for this hypothesis: by expanding \mathcal{P} so that each parameter gets associated a different regularization strength, we can *automatically* identify which parameters are needed and force others towards zero. Fig. 5 illustrates the resulting distribution of per-parameter regularization strengths in the optimal probe for English POS, when \mathcal{P} is defined to be the set of linear probes with per-parameter regularization; interestingly, the dis-




Figure 5: Distribution of the prior precisions learned by linear English POS probes for **Random**, **Word Identity**, **fastText**, and **BERT** representations. High regularization causes weights to be effectively zeroed out.

tribution is bimodal, such that every representation has a set of parameters that is zeroed out (rightmost mode). The random representation is regularized more than pre-trained ones, because it can only learn a majority baseline. Note that in practice, we can do this for probes with multiple layers too, so that the optimal probe we find may be simultaneously deep and sparse.

6 Related Work

Probing aims to provide insights into what linguistic information is encoded in pre-trained representations. Since the introduction of probing for sentence representations (Adi et al., 2017; Conneau et al., 2018), probing has also been applied to representations of words and tokens (Belinkov and Glass, 2019; Liu et al., 2019a; Voita and Titov, 2020; Pimentel et al., 2020b). Nonetheless, comparison of representations, the choice of probe, and even probing tasks have been under scrutiny recently (Belinkov and Glass, 2019; Liu et al., 2019a; Hewitt and Liang, 2019; Pimentel et al., 2020b).

Measuring representation quality. Prior work has mostly used probe accuracy as a measure of the quality of a representation. However, if not properly cross-validated, this can lead to nonsensical results which suggest that random representations are as good as learned ones (Zhang and Bowman, 2018; Hewitt and Liang, 2019). To alleviate this problem, control tasks (Hewitt and Liang, 2019), fewer data (Zhang and Bowman, 2018), or simplistic probes (Liu et al., 2019a) have been used. Using the evidence can be seen as extensive cross-validation (Fong and Holmes, 2020) and is therefore better suited for comparing representations.

In recent work, Lovering et al. (2021) argue that

the ease of extraction of relevant features can be seen as an inductive bias. Specifically, they present experiments on artificial and naturalistic tasks that suggest that the amount of fine-tuning data required to make models rely on relevant features as opposed to spurious correlates of the output is connected to the relative ease of extraction between the spurious and relevant features. In comparison, our method can be seen as integrating over the entire space of features that a representation offers, and as such makes no assumptions about how a task should be solved, i.e., whether certain features are spurious or not for the task at hand.

Simple or complex probes? The choice of probe architecture is still a point of contention in the literature. Initially probes were typically linear models (Alain and Bengio, 2017; Adi et al., 2017; Liu et al., 2019a) because complex probes could memorize and overfit (Zhang and Bowman, 2018; Hewitt and Liang, 2019). However, restricting ourselves to linear probes only allows us to ask whether a particular task has a linear decision boundary, which tells us little about the information encoded in representations. Therefore, neural probes have recently been used as well (Pimentel et al., 2020b; Voita and Titov, 2020). In particular, this has spawned a line of work on automatically trading off probe performance and complexity. For example, Hewitt and Liang (2019) propose control tasks that mitigate overfitting and find that weight decay helps generalization in line with our observations in §5.2. Voita and Titov (2020) use the minimum description length (MDL) principle which is equivalent to the evidence in the case of a probabilistic model (MacKay, 2003). Both of these frameworks focus on the comparison and selection of probes which we argue is distinct from the problem of comparing representations. Thus in our framework, two representations do not need to be compared using the same probe but on the basis of the optimal probe for the representation, which appears to be useful (§5). In this sense, our work is most similar to Pimentel et al. (2020a), where representations, as opposed to probes, are compared by considering the Pareto hypervolume. That said, their approach is dependent on the choice of a complexity metric, whereas ours is not.

Linear probes can overfit. Our results indicate that, for some tasks, even linear probes may be overparameterized. One possible reason for this is that

the optimal probes for these tasks ignore portions of the representation. If true, this would suggest that our framework may be useful for neuron-level probing (Dalvi et al., 2019; Durrani et al., 2020; Torroba Hennigen et al., 2020; Antverg and Belinkov, 2022), whose goal is to identify subsets of neurons in a representation that are informative about a property of interest.

7 Conclusion

Previous approaches to linguistic probing are plagued by several key problems, namely the issues of nonsensical results, probe selection, and ill-defined goals. To overcome these issues, we have proposed a novel probing framework, which focuses on the inductive bias that pre-trained representations offer for different linguistic tasks. We have shown that the Bayesian evidence, a natural measure for inductive bias, can be used in the context of probing. We have found that our framework empirically does not suffer from the aforementioned problems. We are hopeful that under this new paradigm, future work in probing will be more principled, comparable, and useful to the NLP community at large.

Ethics Statement

The authors foresee no ethical concerns with the work presented in this paper.

Acknowledgements

The authors thank Tiago Pimentel, Karolina Stańczak, and members of the McGill NLP group for discussions and providing feedback in various stages of this project, and the anonymous reviewers for their valuable feedback. A. I. acknowledges funding by the Max Planck ETH Center for Learning Systems (CLS). L. T. H. acknowledges funding from the Michael Athans Fellowship fund. V. F. acknowledges funding through a PhD fellowship from the Swiss Data Science Center. R. C. acknowledges support from the Swiss National Science Foundation (SNSF) as part of the “The Forgotten Role of Inductive Bias in Interpretability” project.

References

Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. *Fine-grained analysis of sentence embeddings using auxiliary prediction tasks*. In *International Conference on Learning Representations*.

- Guillaume Alain and Yoshua Bengio. 2017. [Understanding intermediate layers using linear classifier probes](#). In *International Conference on Learning Representations (Workshop Track)*.
- Omer Antverg and Yonatan Belinkov. 2022. [On the pitfalls of analyzing individual neurons in language models](#). In *International Conference on Learning Representations*.
- Roy Bar Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, Bernardo Magnini, and Idan Szpektor. 2006. [The second PASCAL recognising textual entailment challenge](#).
- Oren Barkan. 2017. [Bayesian neural word embedding](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 3135–3143. AAAI Press.
- Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. 2017. [What do neural machine translation models learn about morphology?](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 861–872, Vancouver, Canada. Association for Computational Linguistics.
- Yonatan Belinkov and James Glass. 2019. [Analysis methods in neural language processing: A survey](#). *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Luisa Bentivogli, Ido Dagan, Hoa Trang Dang, Danilo Giampiccolo, and Bernardo Magnini. 2009. [The fifth PASCAL recognising textual entailment challenge](#).
- Christopher M. Bishop. 2006. *Pattern Recognition and Machine Learning*. Springer-Verlag, Berlin, Heidelberg.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. [What you can cram into a single \$\\$&!#*\$ vector: Probing sentence embeddings for linguistic properties](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136, Melbourne, Australia. Association for Computational Linguistics.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. [The PASCAL recognising textual entailment challenge](#). In *Machine Learning Challenges*, pages 177–190. Springer.
- Fahim Dalvi, Nadir Durrani, Hassan Sajjad, Yonatan Belinkov, Anthony Bau, and James Glass. 2019. [What is one grain of sand in the desert? Analyzing individual neurons in deep NLP models](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:6309–6317.
- Francesco D’Angelo and Vincent Fortuin. 2021. [Repulsive deep ensembles are Bayesian](#). In *Advances in Neural Information Processing Systems*. MIT Press.
- Erik Daxberger, Agustinus Kristiadi, Alexander Immer, Runa Eschenhagen, Matthias Bauer, and Philipp Hennig. 2021. [Laplace redux - effortless Bayesian deep learning](#). In *Advances in Neural Information Processing Systems*.
- Marie-Catherine De Marneffe, Mandy Simons, and Judith Tonhauser. 2019. [The CommitmentBank: Investigating projection in naturally occurring discourse](#). In *Proceedings of Sinn und Bedeutung*, volume 23, pages 107–124.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nadir Durrani, Hassan Sajjad, Fahim Dalvi, and Yonatan Belinkov. 2020. [Analyzing individual neurons in pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4865–4880, Online. Association for Computational Linguistics.
- Edwin Fong and Chris C. Holmes. 2020. [On the marginal likelihood and cross-validation](#). *Biometrika*, 107(2):489–496.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. [The third PASCAL recognizing textual entailment challenge](#). In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 1–9, Prague. Association for Computational Linguistics.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. [Learning word vectors for 157 languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).

- John Hewitt and Percy Liang. 2019. [Designing and interpreting probes with control tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, Hong Kong, China. Association for Computational Linguistics.
- Alexander Immer, Matthias Bauer, Vincent Fortuin, Gunnar Rätsch, and Khan Mohammad Emtiyaz. 2021a. [Scalable marginal likelihood estimation for model selection in deep learning](#). In *International Conference on Machine Learning*, pages 4563–4573.
- Alexander Immer, Maciej Korzepa, and Matthias Bauer. 2021b. [Improving predictions of Bayesian neural nets via local linearization](#). In *International Conference on Artificial Intelligence and Statistics*, pages 703–711.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *International Conference on Learning Representations*.
- Christo Kirov, Ryan Cotterell, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sabrina J. Mielke, Arya McCarthy, Sandra Kübler, David Yarowsky, Jason Eisner, and Mans Hulden. 2018. [UniMorph 2.0: Universal Morphology](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [ALBERT: A lite BERT for self-supervised learning of language representations](#). In *International Conference on Learning Representations*.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019a. [Linguistic knowledge and transferability of contextual representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, page 1073–1094. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. [RoBERTa: A robustly optimized BERT pretraining approach](#). *arXiv:1907.11692 [cs]*.
- Charles Lovering, Rohan Jha, Tal Linzen, and Ellie Pavlick. 2021. [Predicting inductive biases of pre-trained models](#). In *International Conference on Learning Representations*.
- David J C MacKay. 1995. [Probable networks and plausible predictions—a review of practical Bayesian methods for supervised neural networks](#). *Network: Computation in Neural Systems*, 6(3):469.
- David J C MacKay. 2003. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press.
- James Martens and Roger Grosse. 2015. [Optimizing neural networks with Kronecker-factored approximate curvature](#). In *International Conference on Machine Learning*, pages 2408–2417.
- Arya D. McCarthy, Miikka Silfverberg, Ryan Cotterell, Mans Hulden, and David Yarowsky. 2018. [Marrying universal dependencies and universal morphology](#). In *Proceedings of the Second Workshop on Universal Dependencies (UDW 2018)*, pages 91–101, Brussels, Belgium. Association for Computational Linguistics.
- Grégoire Montavon, Mikio L. Braun, and Klaus-Robert Müller. 2011. [Kernel analysis of deep networks](#). *Journal of Machine Learning Research*, 12(9).
- Kevin P Murphy. 2012. *Machine learning: A probabilistic perspective*. MIT Press.
- Jerzy Neyman and Egon Sharpe Pearson. 1933. [On the problem of the most efficient tests of statistical hypotheses](#). *Philosophical Transactions of the Royal Society A*, 231(694-706):289–337.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, page 2227–2237. Association for Computational Linguistics.
- Tiago Pimentel and Ryan Cotterell. 2021. [A Bayesian framework for information-theoretic probing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2869–2887, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tiago Pimentel, Naomi Saphra, Adina Williams, and Ryan Cotterell. 2020a. [Pareto probing: Trading off accuracy for complexity](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3138–3153, Online. Association for Computational Linguistics.
- Tiago Pimentel, Josef Valvoda, Rowan Hall Maudslay, Ran Zmigrod, Adina Williams, and Ryan Cotterell. 2020b. [Information-theoretic probing for linguistic structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4609–4622, Online. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.

- Carl Edward Rasmussen and Zoubin Ghahramani. 2000. [Occam’s razor](#). In *Advances in Neural Information Processing Systems*, pages 294–300. MIT Press.
- Lucas Torroba Hennigen, Adina Williams, and Ryan Cotterell. 2020. [Intrinsic probing through dimension selection](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 197–216, Online. Association for Computational Linguistics.
- Luke Vilnis and Andrew McCallum. 2015. [Word representations via gaussian embedding](#). In *International Conference on Learning Representations*.
- Elena Voita and Ivan Titov. 2020. [Information-theoretic probing with minimum description length](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 183–196, Online. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [SuperGLUE: A stickier benchmark for general-purpose language understanding systems](#). In *Advances in Neural Information Processing Systems*, pages 3261–3275.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Boyang Xue, Jianwei Yu, Junhao Xu, Shansong Liu, Shoukang Hu, Zi Ye, Mengzhe Geng, Xunying Liu, and Helen Meng. 2021. [Bayesian transformer language models for speech recognition](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7378–7382. IEEE.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. [XLNet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems*, pages 5754–5764.
- Daniel Zeman, Joakim Nivre, Mitchell Abrams, Noëmi Aepli, Željko Agić, and Lars et al. Ahrenberg. 2019. [Universal dependencies 2.5](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Kelly Zhang and Samuel Bowman. 2018. [Language modeling teaches you more than translation does: Lessons learned through auxiliary syntactic task analysis](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 359–361, Brussels, Belgium. Association for Computational Linguistics.

A Probing as Hypothesis Testing

The result of maximizing over probe P and representation R in eq. (3) can also be understood as selecting the highest-scoring hypothesis in a likelihood-based hypothesis test (Neyman and Pearson, 1933). To see this, we treat the pair (R, P) as one variable M . Then, we compare two representation–probe pairs, M and M' , which we both believe to be equally probable, in the light of data by their likelihood ratio

$$\begin{aligned} \frac{p(M \mid \boldsymbol{\pi}, \boldsymbol{\tau})}{p(M' \mid \boldsymbol{\pi}, \boldsymbol{\tau})} &= \frac{p(\boldsymbol{\pi} \mid \boldsymbol{\tau}, M) p(M) p(\boldsymbol{\pi} \mid \boldsymbol{\tau})}{p(\boldsymbol{\pi} \mid \boldsymbol{\tau}, M') p(M') p(\boldsymbol{\pi} \mid \boldsymbol{\tau})} \\ &= \frac{p(\boldsymbol{\pi} \mid \boldsymbol{\tau}, M)}{p(\boldsymbol{\pi} \mid \boldsymbol{\tau}, M')}. \end{aligned} \quad (5)$$

The simplification is due to the equal probability of M and M' , i.e., $p(M) = p(M')$. If the resulting ratio is larger than one, we decide in favor of M , otherwise in favor of M' . Thus, for a larger set \mathcal{P} of probes and \mathcal{R} of representations, it is sufficient to identify the highest-scoring pair in terms of the evidence $p(\boldsymbol{\pi} \mid \boldsymbol{\tau}, M)$, which is exactly the objective in eq. (3).

B Experimental Details

All our probes are trained using the Adam (Kingma and Ba, 2015) optimizer with hyperparameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, learning rate 0.1, batch size 512, and for 500 epochs. For each discrete architecture (linear, MLP-1, MLP-2), we run the evidence framework as suggested by Immer et al. (2021a) with the following parameters: frequency $F = 1$, $K = 100$ number of steps every epoch, learning rate $\gamma = 0.1$. The evidence framework provides a Laplace approximation to the evidence in eq. (4). We implement our probing method using `laplace-torch` (Daxberger et al., 2021) and use a Kronecker-factored Hessian approximation (Martens and Grosse, 2015) using the eigendecomposition for the Laplace approximation (Immer et al., 2021b). We use an individual weight decay parameter per parameter group of the probes, i.e., each set of weights and biases are regularized independently per layer. Only for Fig. 5, we use a weight decay strength individually per parameter of the linear probe which effectively turns off individual parameters by increasing their weight decay. This is also known as automatic relevance determination (MacKay, 1995). In this case, we use a diagonal Hessian, and thus posterior, approximation.

C Representations

Tab. 1 shows the representations we used. For all transformer models, we use the HuggingFace transformers library (Wolf et al., 2020). Note that for fastText we use the multilingual vectors which are language-dependent, and the official fastText library.⁹

Representation	Model name
m-BERT	bert-base-multilingual-cased
BERT	bert-base-uncased
fastText	Language-specific, see here .
T5	t5-base
RoBERTa	roberta-base
XLNet	xlnet-base-cased
ALBERT	albert-base-v2

Table 1: Representations used. All representation except fastText use the HuggingFace implementations (Wolf et al., 2020).

⁹<https://pypi.org/project/fasttext/>